Very-high-precision solutions of a class of Schrödinger equations

Asif Mushtaq[b], Amna Noreen[a], Kåre Olaussen[a,c], Ingjald Øverbø[a]

[a] *Institutt for fysikk, NTNU*
[b] *Institutt for matematiske fag, NTNU*
[c] *NORDITA, Stockholm*

Abstract

We investigate a method to solve a class of Schrödinger equation eigenvalue problems numerically to very high precision P (from thousands to a million of decimals). The memory requirement, and the number of high precision algebraic operations, of the method scale essentially linearly with P when only eigenvalues are computed. However, since the algorithms for multiplying high precision numbers scale at a rate between $P^{1.6}$ and $P\,\log P\,\log\log P$, the time requirement of our method increases somewhat faster than P^2.

Keywords:
02.30.Hq; 02.30.Mv; 03.65.Ge

1. Introduction

The one-dimensional anharmonic oscillator have been subject to much investigation since the seminal works by Bender and Wu [1, 2] on the behaviour of its perturbation expansion. The motivation has often been to extract features and understanding that can be generalized to more interesting situations, like quantum field theories in higher space-time dimensions, or to test new approximation methods.

In this note we report briefly from our work on a class of one-dimensional quantum mechanical systems which include the one mentioned above, i.e., systems which are modelled by Schrödinger equations of the type

$$-s^2\psi''(x)+\left(x^{2M}+\sum_{m=0}^{M-1}v_m\,x^{2m}\right)\psi(x)=\varepsilon\,\psi(x),\qquad(1)$$

for some finite (small) M, and real coefficients s and v_m. Our focus has been on the possible accuracy to which the eigenvalues and eigenfunctions can be found within available computational resources (memory and CPU cycles). The algorithm we have implemented has modest memory demands; the required memory scales asymptotically with P like MP, where P is the desired precision of eigenvalues or eigenfunctions in decimal digits. The number of required algebraic operations (involving high-precision numbers) generally also seems to grow asymptotically with P like MP. However, in some cases there is a large offset which makes it computationally very expensive to obtain P to even a few digits. Further, the time required per high-precision algebraic operation (i.e. multiplication or division) increases somewhat faster than linearely with P. The high-precision numerical library (CLN [5], built on GMP [6]) we have used has not been parallelized. Thus, our algorithm is mostly constrained by wall-clock time.

Due to space constraints we can in the remainder of this note only present examples of our results (sections 2–4) and a brief decription of the requirements for obtaining a desired precision (section 5). Solving the differential equations (1) numerically to very high precision is the most simple and straightforward part of our work (section 6); the analysis of inevitable loss of accuracy might be more interesting (section 7). The behaviour of our numerical algorithm is illustrated in section 8.

A more complete description will be given elsewhere [7].

2. Ground state energy to one million decimals

As our first *proof-of-method* we considered the ground state of the pure anharmonic oscillator,

$$-\psi''(x)+x^4\,\psi(x)=\varepsilon\,\psi(x),\qquad(2)$$

and computed its ground state energy to $1\,000\,000^+$ decimals. The result, obtained after about 20 days of computing, is

$$\begin{aligned}
&\downarrow\text{Decimal number 1}\\
\varepsilon_0 = 1.&060\,362\,090\,484\,182\,899\,647\,046\,016\,692\,663\backslash\\
&545\,515\,208\,728\,528\,977\,933\,216\,245\,241\,695\backslash\\
&943\,563\,044\,344\,421\,126\,896\,299\,134\,671\,703\backslash\\
&\downarrow\text{Decimal 1 000}\\
\ldots&\,304\,916\,644\,281\,633\,946\,163\,324\,287\,004\,261\backslash\\
&\downarrow\text{Decimal 10 000}\\
\ldots&\,578\,044\,164\,777\,855\,042\,412\,917\,855\,188\,328\backslash\\
&\downarrow\text{Decimal 100 000}\\
\ldots&\,857\,326\,052\,850\,064\,563\,492\,099\,229\,730\,278\backslash\\
&\downarrow\text{Decimal 1 000 000}\\
\ldots&\,820\,139\,466\,721\,621\,064\,477\,821\,481\,635\,914\ldots
\end{aligned}\qquad(3)$$

Email addresses: Asif.Mushtaq@math.ntnu.no (Asif Mushtaq), Amna.Noreen@ntnu.no (Amna Noreen), Kare.Olaussen@ntnu.no (Kåre Olaussen), Ingjald.Overbo@ntnu.no (Ingjald Øverbø)

The full sequence is available on request (for people seeking diversion from investigating the numerical patterns of π).

3. Excited states compared to WKB results

The result of the previous section may be somewhat unsatisfactory to sceptical readers, since there is to our knowledge no similar results to compare against. However, it is only slightly more challenging to treat excited states. To be fair and square we have computed eigenvalue number $n = 50\,000$ of equation (2) to $50\,000^+$ decimals. For such values of n the WKB approximation should be resonably good. Our result is

$$\begin{aligned} &\downarrow \text{Decimal number 1} \\ \varepsilon_{50\,000} = 4\,024\,985.&730\,438\,698\,704\,313\,888\,104\,230\,563\backslash \\ &241\,821\,769\,405\,166\,607\,313\,872\,288\backslash \\ &953\,655\,475\,876\,981\,078\,813\,733\,788\backslash \\ &\downarrow \text{Decimal } 50\,000 \\ \ldots &\,545\,947\,155\,500\,441\,209\ldots \end{aligned} \tag{4}$$

In comparison, the 12^{th} order WKB approximation computed by Bender *et. al.* [4] gives

$$\begin{aligned} \varepsilon_{50\,000}^{\text{WKB-12}} = 4\,024\,985.&730\,438\,698\,704\,313\,888\,104\,230\,563\backslash \\ &241\,821\,769\,405\,166\,607\,313\,872\,288\backslash \\ &953\,657\ldots \end{aligned} \tag{5}$$

I.e., the relative accuracy of the WKB approximation is

$$\frac{\varepsilon_{50\,000}^{\text{WKB-12}} - \varepsilon_{50\,000}}{\varepsilon_{50\,000}} = 5.163\ldots \times 10^{-67}. \tag{6}$$

From observation of the behaviour of the WKB series we find this accuracy to be as expected for the 12^{th} order approximation at this value of n. Some of us plan to return to a more detailed analysis of the behaviour of the WKB approximation, which may nowadays be extended easily to much higher orders.

4. Brute force calculation of double-well level-splitting

Another well analysed situation where we may stress-test our method is the calculation of the level splitting between the lowest even and odd parity eigenstates of the double-well potential,

$$-s^2\psi''(x) + \left(x^2-1\right)^2\psi(x) = \varepsilon\,\psi(x), \tag{7}$$

for small s. The lowest even, $\varepsilon_0^{(+)}$, and odd, $\varepsilon_0^{(-)}$, parity states are split by an exponentially small amount $\Delta\varepsilon_0 \equiv \varepsilon_0^{(-)} - \varepsilon_0^{(+)}$. Asymptotically as $s \to 0^+$,

$$\Delta\varepsilon_0 \sim \Delta\varepsilon_0^{\text{Z-J}} = 16\sqrt{\frac{2s}{\pi}}\,\mathrm{e}^{-4/3s}\,\mathrm{e}^{L(s)}, \tag{8}$$

where $L(s) = -\left(\frac{71}{96}s + \cdots\right)$ is given to order s^{10} by Zinn-Justin[1] [3]. We have made independent calculations of $\varepsilon_0^{(\pm)}$ to $30\,000^+$ digits accuracy for $s = \frac{1}{50\,000}$:

$$\begin{aligned} &\text{Decimal } 28\,954 \downarrow \\ \varepsilon_0^{(-)} &= 0.000\,039\,999\,799\ldots 990\,905\,404\ldots \end{aligned} \tag{9}$$

$$\varepsilon_0^{(+)} = 0.000\,039\,999\,799\ldots 984\,723\,697\ldots \tag{10}$$

Equation (8) agrees with the difference to the expected order,

$$\frac{\Delta\varepsilon_0^{\text{Z-J}} - \Delta\varepsilon_0}{\Delta\varepsilon_0} = 1.649\ldots \times 10^{-48} \approx 8\,052\,s^{11}. \tag{11}$$

The right hand side is of the magnitude expected for the next term in $L(s)$.

We hope the three examples above have convinced the reader that it is possible to solve the eigenvalue problems (1) to very high precision. *How* precise will of course depend on the parameters and which eigenstate we want to investigate.

5. How to achieve a desired precision

The eigenvalue condition for equation (1) is assumed to be that $\psi(x) \to 0$ as $x \to \pm\infty$. We are unable to impose this exactly in our numerical algorithm. However, there is an equivalent Robin boundary condition which can be imposed at some finite (large) x,

$$-s\frac{\psi'(x)}{\psi(x)} = R(x) = x^M + \cdots \approx \infty. \tag{12}$$

We don't know $R(x)$ exactly, but there is for any desired precision P a finite value of x such that an approximate $R(x)$ is sufficient. The required value of x can be estimated by asymptotic analysis of equation (1) as $x \to \infty$ (or a WKB approximation to include estimates of constant prefactors which cannot be found by asymptotic analysis alone). A first estimate is that one should choose x so that

$$\exp\left(-\frac{2}{s}\int_{x_0}^{x}\sqrt{V(y)-\varepsilon}\,\mathrm{d}y\right) \equiv 10^{-P_{\text{est}}(x)} \lesssim 10^{-P}, \tag{13}$$

if one wants to compute the eigenvalue to P decimals precision. Here x_0 is the largest turning point, and to simplify we have assumed that the Robin boundary condition is replaced by a Diriclet one, $R(x) = \infty$.

Equation (13) is an *a priori* estimate, which we have tested by choosing a very large x to obtain a very accurate eigenvalue (so that it may be considered exact), and used this to observe the obtainable precision for lower values of x. The obtainable precision at a given x is defined as

$$P(x) \equiv \lg|\varepsilon(x) - \varepsilon|\,, \tag{14}$$

where $\varepsilon(x)$ is the eigenvalue found numerically by use of equation (12). Figure 1 displays the difference between $P_{\text{est}}(x)$ and $P(x)$ in some cases. Note that the difference between the estimated $P_{\text{est}}(x)$ and obtainable $P(x)$ precision varies very little with x, and hence can be found numerically from fairly low-precision calculations. The difference probably occurs because we have neglected a slowly varying prefactor in (13).

[1] We use a different normalization: $s = 8g$ and $\varepsilon = 32gE$, where g and E are the parameters in [3].

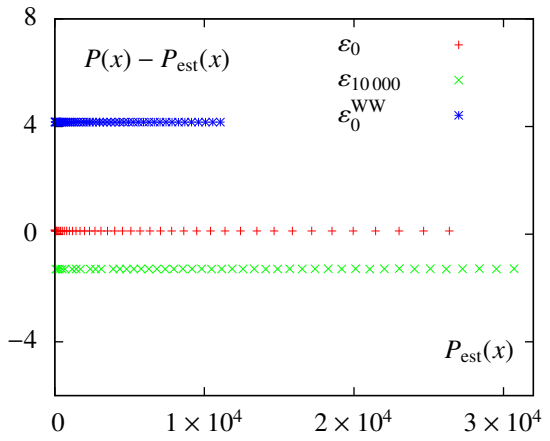

Figure 1: The obtainable precision $P(x)$ for various eigenvalues ε with a Diriclet boundary condition (12) imposed at x instead of the exact condition. The cases plotted are (i) for the lowest eigenvalue ε_0 of equation (2), for which $P_{\rm est}(x) \approx 2x^3/3\ln 10$, (ii) for the $10\,000^{\rm th}$ eigenvalue $\varepsilon_{10\,000}$, and (iii) for the lowest eigenvalue $\varepsilon_0^{\rm WW}$ of equation (7), for which $P_{\rm est}(x) \approx 2(x-1)^2(x+2)/(3s\ln 10)$.

It is reassuring that the obtainable precision can be predicted to within a few digits: Choosing a too small x leads to the solution of the wrong eigenvalue problem; choosing a too large x leads to a waste of CPU cycles. The obtainable precision at fixed x may be improved by some number of decimals by using a better Robin boundary condition found by asymptotic analysis of equation (1). Such improvement might be useful in situations where a few tens of decimals precision is sufficient.

6. Method for solving the Schrödinger equation

We have postponed the description of our numerical method of solving equation (1), due to its (perhaps) disappointingly naive simplicity: We make a brute force summation of its Taylor expansion,

$$\psi(x) = x^\sigma \sum_{m=0}^{N} a_m x^{2m}, \tag{15}$$

where $\sigma = 0$ or 1 depending on the parity of the solution. The coefficients a_m, or more precisely $A_m(x) \equiv a_m\, x^{2m}$, are generated recursively from equation (1), starting with $A_0(x) = 1$. Only the $M+1$ last coefficients need to be considered at any time while the sum is accumulated; hence the memory requirement is low.

Since equation (1) has no singular points in the finite plane the sum will eventually converge *very fast*. The number of terms N needed in the sum (15) can be chosen automatically by the recursion/summation routine, but may also be *a priori* estimated. The calculation is done in very-high-precision floating-point arithmetic using the CLN C++ library of numbers [5].

7. Numerical loss of accuracy

Note that for the harmonic oscillator the summation (15) means computing its ground state $\psi_0(x) = \mathrm{e}^{-x^2/2}$ for large x by Taylor expansion. This is certainly not the recommended method of computation, due to large cancellations and roundoff errors. However, by calculating with sufficient numerical precision — which is not prohibitively large — it actually works quite well. Further, the location of eigenvalues are determined by ε-values where $\psi(x) = \psi(x;\varepsilon)$ changes sign very rapidly with ε. The computation of eigenvalues only is less sensitive to cancellations.

Nevertheless, the effects of roundoff and cancellations must be considered. Computing with high-precision floating point numbers with D decimals accuracy means that the value of $\psi(x)$ in equation (15) is accumulated from numbers with a mantissa of $(D\ln 10/\ln 2)$ bits. A contribution of magnitude $10^{\Delta D}$ to the sum will thus have a round-off error of order $10^{\Delta D - D}$. S terms of the same magnitude is expected to increase this error by a factor $\sqrt{S}$ (with symmetric roundoff), which is an insignificant increase. In principle there might also be error amplification in the recursion relation, but we have not observed signatures of such. Thus, we estimate the numerical accuracy loss to ΔD decimal digits, where

$$10^{\Delta D} = \max_m \{\, |A_m(x)| \,\}. \tag{16}$$

The largest term in the sum (15) is found numerically by monitoring the recursion/summation routine. It may also be *a priori* estimated. We have done the latter in two ways: First, by asymptotic analysis of the recursion relations in some simple situations (those described in sections 2 and 4). This analysis rapidly becomes complicated. Second, all analytic and numerical results found are consistent with the assumption that

$$\max_m \{\, |A_m(x)| \,\} = \max_\varphi \left\{\, |\psi(x\mathrm{e}^{\mathrm{i}\varphi})| \,\right\}, \tag{17}$$

where $\psi(x\mathrm{e}^{\mathrm{i}\varphi})$ can be estimated from a WKB-approximation. Equation (17) is based on the assumptions that (i) there is always a point on the circle $x\mathrm{e}^{\mathrm{i}\varphi}$ where cancellations are insignificant in the sum (15), and (ii) the main contributions to the sum come from relatively few terms around the maximum term.[2] There may be parameter combinations where the first assumption fails, in which case the equality sign in equation (17) should be replaced by $\geq$ (less helpful for estimations).

For the example in section 2 we find

$$P_{\rm est}(x) = \frac{2\,x^3}{3\ln 10}, \quad \Delta D = \frac{x^3}{3\ln 10}. \tag{18}$$

Thus, to compute the ground state to P decimals accuracy we must evaluate the wavefunction at $x = \left[\left(\frac{3}{2}\ln 10\right)P\right]^{1/3}$. If we

[2] Actually, the number of terms contributing to $\max_\varphi \left\{\, |\psi(x\mathrm{e}^{\mathrm{i}\varphi})| \,\right\}$ is an unimportant correction to ΔD. For the example in section 2 we used $x = 152$, for which the estimated maximum is about $10^{508\,386}$, and summed less than 10^7 terms of the Taylor series. Whether the maximum is contributed from one single or almost all terms of the sum makes only a few decimals change in ΔD.

want to evaluate the wavefunction to P decimals at this x we must choose a numerical precision of $D = \frac{3}{2}P$ decimals, since we will loose $\Delta D = \frac{1}{2}P$ decimals to roundoff errors. However, if we only want to find the eigenvalue ε we experience a compensating *accuracy gain*. An uncertainty $\delta\psi$ in the wavefunction translates to an uncertainty

$$\delta\varepsilon = \left(\frac{\partial\psi}{\partial\varepsilon}\right)^{-1} \delta\psi \tag{19}$$

in the eigenenergy. In this example $(\partial\psi/\partial\varepsilon) \approx \mathrm{e}^{x^3/3} \approx 10^{P/2}$. Hence, the accuracy loss due to roundoff is completely compensated by the accuracy gain caused by a very large $(\partial\psi/\partial\varepsilon)$. This implies that our method of computing eigenvalues may work well even for ordinary precisions P. There are many cases where such complete compensation occur.

The example in section 4 is different. We find

$$\begin{aligned} P_{\mathrm{est}}(x) &\approx \frac{2}{3s\ln 10}\,(x-1)^2\,(x+2)\,, \\ \Delta D &\approx \frac{1}{s\ln 10}\left(\frac{1}{3}x^3 + \frac{1}{2}x\right), \\ \frac{\partial\psi}{\partial\varepsilon} &\approx 10^{P_{\mathrm{est}}(x)/2}. \end{aligned} \tag{20}$$

In this case the accuracy loss is not fully compensated. We note that $\Delta D \approx 5/(6s\ln 10)$ (large when s is small) when P is chosen small. In this case it will be computationally very expensive to obtain results to a few decimals of accuracy by this method. However, for very large P the computational cost is similar to the example in section 2. The example in section 3 is similar to the one section 4, only with more cumbersome expressions.

8. Numerical observations

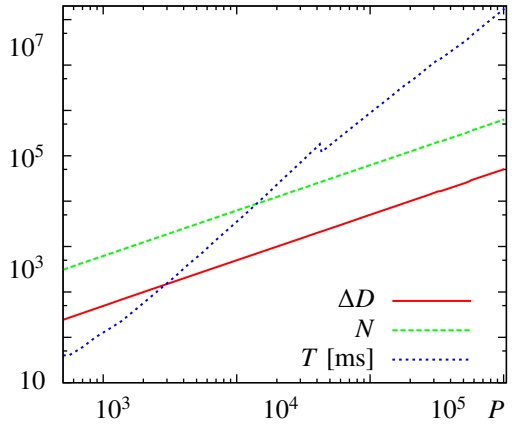

Figure 2: Some observed behaviour when calculating the ground state energy of equation (2) to P decimals precision: (i) The accuracy loss ΔD (in decimals) due to roundoff error; in this case $\Delta D \sim P/2$. (ii) The number N of terms needed in the Taylor expansion (15), proportional to P. (iii) The time T used to evaluate one wavefunction to required precision, locally we find $T \sim P^\nu$, where $\nu \approx 2.6$ around $P = 10\,000$ and $\nu \approx 2.13$ around $P = 200\,000$. (Datapoints for $P > 300\,000$ are not generated under identical conditions.)

We monitor many parameters during the numerical computations. A sample of those are shown in figure 2, from computation of the ground state energy of equation (2). We plot the numerically observed accuracy loss ΔD; this agrees with the *a priori* estimate. This is also observed for all other investigated cases.

Further, the number N of terms required in the sum (15) seems to grow linearly with precision P. This is also the case for other examples, only with different coefficients of proportionality.

The total time to make one evaluation of the wavefunction $\psi(x)$ also behaves as expected from the number of terms in the sum, and the time needed to multiply two high-precision numbers. The drop in computation time near $P = 42\,000$ does not seem to be an artifact of variations in the computational environment. We believe it is due to a change of multiplication algorithm at this precision.

9. Possible extensions

The possibilites of extending our method to other systems are somewhat limited. It seems straightforward to generalize to non-symmetric one-dimensional potentials, to Schrödinger equations which have only one regular singular point in the finite plane, and to small systems of such equations. We believe that systems with two (unseparable) degrees of freedom can be constructively approached.

The evaluation of unnormalized wavefunctions is certainly possible, with a time requirement proportional to the number of evaluation points. We do not rule out the possibility of computing normalized wave functions to high precision.

Another interesting extension is towards very-high-precision computation of Green functions for the same class of models.

10. Acknowledgement

This work was supported in part by the Higher Education Commision of Pakistan (HEC) and the "SMÅFORSK" program of the Research Council of Norway. KO thanks NORDITA, Stockholm for hospitality and support July-December 2009, where parts of this work was performed.

References

[1] C.M. Bender and T.T. Wu, *Large-Order Behaviour of Perturbation Theory*, Physical Review Letters **27**, 461 (1971)
[2] C.M. Bender and T.T. Wu, *Anharmonic Oscillator. II. A Study of Perturbation Theory in Large Order*, Physical Review **D7**, 1620 (1972)
[3] J. Zinn-Justin, *Expansion around instantons in quantum mechanics*, J. Math Phys. **22**, 511 (1981)
[4] C.M. Bender, K. Olaussen and P.S. Wang, *Numerological analysis of the WKB approximation in large order*, Physical Review **D16**, 1740 (1977)
[5] B. Haible and R.B. Kreckel, *CLN – Class Library for Numbers*, http://www.ginac.de/CLN/
[6] T. Granlund and collaborators, *GMP – The GNU Multiple Precision Arithmetic Library*, http://gmplib.org/
[7] A. Mushtaq, A. Noreen, K. Olaussen, and I. Øverbø, In preparation.